

REMARKS

After entry of the above amendments, claims 34-49 will be pending in the present application.

Claims 1-33 have been cancelled. New claims 34-49 have been added. Support for the new claims can be found in the specification, drawings, and claims as originally filed. No new matter has been added.

In this Amendment, Applicant has cancelled previously pending claims 1-33 from further consideration in this application. Applicant is not conceding that the subject matter encompassed by claims 1-33 is not patentable over art cited by the Examiner. Claims 1-33 have been cancelled in this Amendment solely to facilitate expeditious prosecution of the present application. Applicant reserves the right to pursue claims directed to the subject matter encompassed by claims 1-33 and any additional claims in one or more continuing and/or divisional applications.

Claim Objections

Previously pending claims 1-12 and 18 were objected to on the basis of informalities. Since claims 1-12 and 18 have been cancelled and none of the newly added claims recite the informalities noted by the Examiner, Applicant respectfully requests withdrawal of the claim objections.

§ 101 Rejections

Previously pending claims 14-23 were rejected under 35 U.S.C. § 101 as being directed to non-statutory subject matter. Since claims 14-23 have been cancelled and none of the newly added claims recite language noted by the Examiner, Applicant respectfully requests withdrawal of the claim rejections under § 101.

Double Patenting

Previously pending claims 1-33 were provisionally rejected under 35 U.S.C. § 101 as claiming the same invention as that of claims 1-33 of co-pending U.S. Patent Application No. 10/763,135. Applicant has cancelled claims 1-33 from the present application and claims 1-33 from co-pending U.S. Patent Application No. 10/763,135. Applicant respectfully submits that the newly added claims in the present application and the newly added claims in co-pending U.S. Patent Application No. 10/763,135 are patentably distinct as the newly added claims in the present application recite elements not found in the newly added claims of co-pending U.S. Patent Application No. 10/763,135, and vice versa. Therefore, withdrawal of the provisional double patenting rejection is respectfully requested.

§ 102 / § 103 Rejections

Previously pending claims 1, 5-7, 11-14, 17-19, 22-24, and 28-32 were rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 7,051,098 to Masters et al. (hereinafter “Masters”). Previously pending claims 8-10, 20-21, and 33 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Masters, in view of U.S. Patent Application Publication No. 2003/0208523 to Gopalan et al. (hereinafter “Gopalan”). Previously pending claims 2-4, 15-16, and 25-27 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Masters, in view of U.S. Patent No. 5,841,869 to Merkling et al. (hereinafter “Merkling”).

New claim 34 recites:

34. A method for supporting a transaction application workload and a parallel application workload on one server cluster, the method comprising:
 receiving a request from a client to execute the transaction application workload on the one server cluster, the one server cluster including server nodes at one domain;
 identifying a service level agreement negotiated with the client for the transaction application workload, the service level agreement specifying performance

requirements for execution of the transaction application workload on the one server cluster;

assigning a subset of the server nodes in the one server cluster at the one domain to execute the transaction application workload;

monitoring execution of the transaction application workload on the subset of server nodes assigned to execute the transaction application workload to determine whether the performance requirements for execution of the transaction application workload specified in the service level agreement are being met; and

responsive to a determination that the performance requirements for execution of the transaction application workload specified in the service level agreement are not being met, dynamically reassigning one or more of the server nodes in the one server cluster at the one domain assigned to execute the parallel application workload to the execution of the transaction application workload in order to meet the performance requirements for execution of the transaction application workload specified in the service level agreement,

wherein a server node assigned to execute the transaction application workload cannot be concurrently assigned to execute the parallel application workload and a server node assigned to execute the parallel application workload cannot be concurrently assigned to execute the transaction application workload.

In the Office action, the Examiner states:

Masters teaches . . . in response to the identified violation, dynamically reallocating a computing resource assigned to the parallel application to the transaction application that requires an additional computing resource to meet the service level agreement (when an application is not meeting its performance requirements the QoS managers will request the resource manager FG42 scale up a new copy or move the application to a new host, see, e.g., col. 37, lines 47-58) . . .

(October 11, 2007 Office action, pgs. 3-4).

Instantiating a new copy of an application or moving an application from one host to another host CANNOT be construed as disclosing “dynamically reassigning one or more of the server nodes in the one server cluster at the one domain assigned to execute the parallel application workload to the execution of the transaction application workload”, as recited in claim 34.

To give an example, suppose server nodes A and B are assigned to execute a transaction application workload and server nodes C and D are assigned to execute a parallel application workload. “Dynamically reassigning one or more of the server nodes . . . assigned to execute the parallel

application workload to the execution of the transaction application workload" would involve reassigning server node C, server node D, or both server nodes C and D to execute the transaction application workload. In other words, additional server node(s) are assigned to the execution of the transaction application workload.

Movement of an application from one host to another host definitely does NOT involve assigning additional host(s) to execution of an application since one host is simply substituting for another host. In addition, it is NOT inherent that instantiation of a new copy of an application involves instantiating the new copy on an additional host. Under MPEP § 2112:

"To establish inherency, the extrinsic evidence 'must make clear that the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill. Inherency, however, may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient.' " *In re Robertson*, 169 F.3d 743, 745, 49 USPQ2d 1949, 1950-51 (Fed. Cir. 1999) (citations omitted).

(MPEP § 2112, 8th ed., Sept. 2007 rev.).

Persons of ordinary skill in the art readily understand that a new copy of an application can be instantiated on the same host and need not be instantiated on a different host. Therefore, it is NOT inherent that the instantiation of a new copy of an application in Masters is on an additional host.

Even assuming argumentatively that instantiation of a new copy of an application in Masters is on an additional host, claim 34 recites "dynamically reassigning one or more of the server nodes . . . assigned to execute the parallel application workload to the execution of the transaction application workload . . . , wherein a server node assigned to execute the transaction application workload cannot be concurrently assigned to execute the parallel application workload and a server node assigned to execute the parallel application workload cannot be concurrently assigned to execute the transaction

application workload". There is no discussion in Masters about a host only being able to execute one type of application at any one time.

The Examiner cites the following passage of Merkling for teaching parallel and transaction applications:

New generations of kernels provide specific support for adaptability, configurability of their basic components, and the supply of required dedicated run-time behaviors to critical applications, collectively referred to herein as personalities. Applications use various kernel resources, e.g. scheduling components, synchronization components, threads, memory handlers, exception handlers, and interrupts. Earlier research has demonstrated that heavy transactions applications, as well as parallel applications, often require dynamic process migration and load balancing. To meet the fundamental properties required for such resource management, components of constituent products require standard interfaces, data formats, and protocols.

(Col. 12, Ins. 19-32 of Merkling).

The passage of Merkling cited by the Examiner, however, only discusses managing allocation of kernel resources to transaction applications or parallel applications. Allocation of kernel resources is NOT the same as dynamically allocating server nodes in a server cluster. In addition, as with Masters, Merkling does NOT mention anything about a host only being able to execute one type of application at any one time. Therefore, neither Masters nor Merkling deals with the specific problem addressed by the present application.

Although Masters discusses "specify[ing] quality of service (QoS) requirements for each host, each application, and the network in which the hosts are connected" (col. 2, Ins. 29-31 of Masters), claim 34 recites "identifying a service level agreement negotiated with the client for the transaction application workload". As discussed on page 9, paragraph [0029] of the present application, the service level agreements are defined on a workload basis. There is no requirement that workloads of the same type of application have the same service level agreement. In fact, page 10, paragraph [0031] of the present application specifically discusses different classes of workloads in an application and different

service level agreements for the different classes of workloads. Hence, specifying quality of service for an application is NOT the same as “identifying a service level agreement negotiated with the client for the transaction application workload”, as recited in claim 34.

Gopalan does not cure the deficiencies of Masters or Merkling. Therefore, even if Gopalan were combined with Masters and Merkling, the combination would neither teach nor suggest all of the elements of claim 34.

Based at least on the reasons above, Applicant respectfully submits that claim 34, and the claims that depend therefrom, are patentable over Masters, in view of Gopalan, and further in view of Merkling.

New claim 35, which depends upon claim 34, recites:

35. The method of claim 34, further comprising:
 monitoring execution of the transaction application workload on the subset of server nodes assigned to execute the transaction application workload to determine whether the performance requirements for execution of the transaction application workload specified in the service level agreement will continue to be met; and
 responsive to a determination that the performance requirements for execution of the transaction application workload specified in the service level agreement will not continue to be met, dynamically reassigning one or more of the server nodes in the one server cluster at the one domain assigned to execute the parallel application workload to the execution of the transaction application workload in order to continue to meet the performance requirements for execution of the transaction application workload specified in the service level agreement.

In the Office action, the Examiner states:

Masters teaches . . . history servers (40 in figure 1A) is used to drive line graph charts for specific hosts selected at the host display in order to provide CPU load information, network load information and memory utilization information (see, e.g., col. 46, lines 40-61).

Even though Masters teaches implicitly the claim limitation, Gopalan explicitly teaches as follows:

a predictive real-time Service Level Agreements (SLAs) monitoring system by a detailed analysis of traffic flows with reduced monitor data flow across network;

a suitable combination of offline and real-time processing of historical and current traffic data;

the system analyzes the historical traffic patterns of a network to determine a set of critical SLAs; and

the system analyzes the gathered data from the master and slave probes to forecast future violations of the SLA and generate operator SLF violation alarms (see, e.g., abstract).

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Masters to include predictive real-time SLA monitoring system as taught by Gopalan in order to efficiently analyze the historical traffic patterns, forecast future violations of the SLAs and generate operator SLA violation alarms.

(October 11, 2007 Office action, pg. 8).

Gopalan, however, deals with monitoring and predicting network traffic, NOT execution of an application workload on server nodes of a server cluster. Therefore, modifying Masters in view of Gopalan would only involve forecasting future violations of the quality of service specified for networks. It would NOT extend to forecasting future violations of the quality of service specified for hosts or applications because Gopalan does NOT provide any teachings or suggestions concerning the quality of service for hosts or applications.

Since Merkling does NOT cure the deficiencies of Masters and Gopalan, even if Merkling were combined with Masters and Gopalan, the combination would neither teach nor suggest all of the elements of claim 35.

Based at least on the above, Applicant respectfully submits that claim 35 is further patentable over Masters, in view of Gopalan, and further in view of Merkling.

CONCLUSION

On the basis of the above remarks, reconsideration and allowance of the claims is believed to be warranted and such action is respectfully requested. If the Examiner has any questions or comments, the Examiner is respectfully requested to contact the undersigned at the number listed below.

Respectfully submitted,
SAWYER LAW GROUP LLP

Dated: March 27, 2008

/Erin C. Ming/
Erin C. Ming
Attorney for Applicant
Reg. No. 47,797
(650) 475-1449